The I/O complexity of Strassen's matrix multiplication with recomputation

Gianfranco Bilardi[1] and Lorenzo De Stefani[2]

1 **Department of Information Engineering, University of Padova, Via Gradenigo 6B/Padova, Italy**
bilardi@dei.unipd.it

2 **Department of Computer Science, Brown University, 115 Waterman Street/Providence, United States of America**
lorenzo@cs.brown.edu

arXiv:1605.02224v1 [cs.DS] 7 May 2016

Abstract

A tight $\Omega((n/\sqrt{M})^{\log_2 7}M)$ lower bound is derived on the I/O complexity of Strassen's algorithm to multiply two $n \times n$ matrices, in a two-level storage hierarchy with M words of fast memory. A proof technique is introduced, which exploits the Grigoriev's flow of the matrix multiplication function as well as some combinatorial properties of the Strassen computational directed acyclic graph (CDAG). Applications to parallel computation are also developed. The result generalizes a similar bound previously obtained under the constraint of no-recomputation, that is, that intermediate results cannot be computed more than once. For this restricted case, another lower bound technique is presented, which leads to a simpler analysis of the I/O complexity of Strassen's algorithm and can be readily extended to other "Strassen-like" algorithms.

1998 ACM Subject Classification F.2.1 Numerical Algorithms and Problems

Keywords and phrases I/O complexity, Strassen's Matrix Multiplication, Recomputation, Memory Hierarchy, Parallel Computation

1 Introduction

Data movement is increasingly playing a major role in the performance of computing systems, in terms of both time and energy. This current technological trend [19] is destined to continue, since the very fundamental physical limitations on minimum device size and maximum message speed lead to inherent costs when moving data, whether across the levels of a hierarchical memory system or between processing elements of a parallel systems [8]. The communication requirements of algorithms have been the target of considerable research in the last four decades; however, obtaining significant lower bounds based on such requirements remains and important and challenging task.

In this paper we focus on the I/O complexity of Strassen's matrix multiplication algorithm. Matrix multiplication is a pervasive primitive utilized in many applications. Strassen [24] showed that two $n \times n$ matrices can be multiplied with $O(n^\omega)$, with $\omega = \log_2 7 \approx 2.8074$, hence with asymptotically fewer than the n^3 arithmetic operations required by the straightforward implementation of the definition of matrix multiplication. This result has motivated a number of efforts which have lead to increasingly faster algorithms, at least asymptotically, with the current record being at $\omega < 2.3728639$ [17].

© Gianfranco Bilardi and Lorenzo De Stefani;
licensed under Creative Commons License CC-BY
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Previous and related work

I/O complexity has been introduced in the seminal work by Hong and Kung [15]; it is essentially the number of data transfers between the two levels of a memory hierarchy with a fast memory of M words and a slow memory with an unbounded number of words. They presented techniques to develop lower bounds to the I/O complexity of computations modeled by *computational directed acyclic graphs* (CDAGs). The resulting lower bounds apply to all the schedules of the given CDAG, including those with recomputation, that is, where some vertices of the CDAG are evaluated multiple times. Among other results, they established an $\Omega\left(n^3/\sqrt{M}\right)$ lower bound to the I/O complexity of the definition-based matrix multiplication algorithm, which matched a known upper bound upper bound [10]. The techniques of [15] have also been extended to obtain tight communication bounds for the definition-based matrix multiplication in some parallel settings [14, 12].

Ballard et al. generalized the results on matrix multiplication of Hong and Kung [15] in [4, 3] by using the approach proposed in [14] based on the Loomis-Whitney geometric theorem [18, 26]. The same papers presents tight I/O complexity bounds for various classical linear algebra algorithms, for problems such as LU/Cholesky/LDLT/QR factorization and eigenvalues and singular values computation.

It is natural to wonder what is the impact of Strassen's reduction of the number of arithmetic operations on the number of data transfers. In an important contribution Ballard et al. [5], obtained an $\Omega((n/\sqrt{M})^{\log_2 7}M)$ I/O lower bound for Strassen's algorithm, using the "*edge expansion approach*". The authors extend their technique to a class of "*Strassen-like*" fast multiplication algorithms and to fast recursive multiplication algorithms for rectangular matrices [2]. This result was later generalized to a broader class of "*Strassen-like*" algorithms of by Scott et. al [23] using the "*path routing*" technique. A parallel, "*communication avoiding*" implementation of Strassen's algorithm whose performance matches the known lower bound [5, 23], was proposed by Ballard et al. [1].

The edge expansion technique of [5], the path routing technique of [23], and the "*closed dichotomy width*" technique of [9] all yield I/O lower bounds that apply only to computational schedules for which no intermediate result is ever computed more than once (*nr-computations*). While it is of interest to know what is the I/O complexity achievable by nr-computations, it is also important to investigate what can be achieved with recomputation. In fact, for some CDAGs, recomputing intermediate values does reduce the space and/or the I/O complexity of an algorithm [21]. In [6], it is shown that some algorithms admit a *portable schedule* (i.e., a schedule which achieves optimal performance across memory hierarchies with different access costs) only if recomputation is allowed. A number of lower bound techniques that allow for recomputation have been presented in the literature, including the "*S-partition* technique" [15], the "*S-span* technique" [21], and the "*S-covering* technique" [7] which merges and extends aspects from both [15] and [21]. None of these has however been successfully applied to fast matrix multiplication algorithms.

Our results

Our main result is the extension of the $\Omega((n/\sqrt{M})^{\log_2 7}M)$ I/O complexity lower bound for Strassen's algorithm to schedules with recomputation. A matching upper bound is known, and obtained without recomputation; hence, we can conclude that, for Strassen's algorithm, recomputation does not help in reducing I/O complexity if not, possibly, by a constant factor. In addition to the result itself, the proof technique appears to be of independent interest, since it exploits to a significant extent the divide&conquer nature which is exhibited by many

algorithms. We do follow the dominator-set approach pioneered by Hong and Kung in [15]. However, we focus the dominator analysis only on a select set of target vertices, specifically the outputs of the sub-CDAGs of Strassen's CDAG that correspond to sub-problems of a suitable size (a size chosen as a function of the fast memory capacity, M). Any dominator set of a set of target vertices can be partitioned into two subsets, one internal and one external to the sub-CDAGs. The analysis of the external component of the dominator does require rather elaborate arguments that are specific to Strassen's CDAG. In contrast, the analysis of the internal component can be carried out based only on the fact that the sub-CDAGs compute matrix products, irrespective of the algorithm (in our case, Strassen's) by which the products are computed. To achieve this independence of the algorithm, we resort on the concept of Grigoriev's flow of a function [13] and on a lower bound to such flow established by Savage [22] for matrix multiplication.

As it turns out, for schedules without recomputation, the analysis of the external component of the dominator sets becomes unnecessary and the analysis of the internal components can be somewhat simplified. The result is a derivation of the I/O complexity without recomputation considerably simpler than those in [5, 23]. The technique is easily extended to a class of Strassen-like algorithms.

The paper is organized as follows: in Section 2, we provide the details of our model and of several theoretical notions needed in our analysis. Section 3 describes the simplified lower bound for Strassen's and Strassen-like algorithms, when implemented with no recomputation. In Section 4, we present the I/O complexity lower bound for Strassen's algorithm, when recomputation is allowed. Extensions of the result to a parallel model are also discussed.

2 Preliminaries

We consider algorithms which compute the product $C = AB$ of $n \times n$ matrices A, B with entries form a ring $\mathcal{R}$. Specifically, we focus on algorithms whose execution, for any given n, can be modeled as a *computational directed acyclic graph* (CDAG) $G(V, E)$, where each vertex $v \in V$ represents either an input value or the result of an unit time operation (i.e. an intermediate result or one of the output values), while the directed edges in E represent data dependences. A *directed path* connecting vertices $u, v \in V$ is an ordered sequence of vertices for which u and v are respectively the first, and last vertex such that there is in E a (directed) edge pointing from each vertex in the sequence to its successor. We say that a CDAG $G'(V', E')$ is a *sub-CDAG* of $G(V, E)$ if $V' \subseteq V$ and $E' \subseteq E$.

Model

We assume that sequential computations are executed on a system with a two-level memory hierarchy consisting of a fast memory or *cache* of size M, measured in words, and a *slow memory* of unlimited size. We assume that each memory word can store at most one value form $\mathcal{R}$. An operation can be executed only if all its operands are presents in cache. Data can be moved from the slow memory to the cache by *read* operations and in the other direction by *write* operations. Read and write operations are also called *I/O operations*. We assume the input data to be stored in slow memory at the beginning of the computation. The evaluation of a CDAG in this model can be analyzed by means of the "*red-blue pebble game*" [15]. The number of I/O operations executed when evaluating a CDAG depends on the "*computational schedule*," that is, on the order in which vertices are evaluated and on which values are kept in/discarded from cache. The *I/O complexity* $IO_G(M)$ of CDAG G is defined as the minimum number of I/O operations over all possible computational schedules.

We also consider a parallel model where P processors, each with a local memory of size M, are connected by a network. We assume that the input is initially distributed among the processors, thus requiring that $MP \geq 2n^2$. Processors can exchange point-to-point messages among each other. For this model, we derive lower bounds to the number of words that must be either sent or received by at least one processor during the CDAG evaluation.

Preliminary definitions

The concept of *information flow of a function* was originally introduced by Grigoriev [13]. We use a revised formulation presented by Savage [22]. We remark that the flow is an inherent property of a function, not of a specific algorithm by which the function may be computed.

▶ **Definition 1** (Grigoriev's flow of a function)**.** A function $f : \mathcal{R}^p \to \mathcal{R}^q$ has a $w(u, v)$ Grigoriev's flow if for all subsets X_1 and Y_1 of its p input and q output variables, with $|X_1| \geq u$ and $|Y_1| \geq v$, there is a sub-function h of f obtained by making some assignment to variables of f not in X_1 and discarding output variables not in Y_1 such that h has at least $|\mathcal{R}|^{w(u,v)}$ points in the image of its domain.

A lower bound on the Grigoriev's flow for the square matrix multiplication function $f_{n\times n} : \mathcal{R}^{2n^2} \to \mathcal{R}^{n^2}$ over the ring $\mathcal{R}$ was presented by Savage in [22] (Theorem 10.5.1).

▶ **Lemma 2** (Grigoriev's flow of $f_{n\times n} : \mathcal{R}^{2n^2} \to \mathcal{R}^{n^2}$ [22])**.** *$f_{n\times n} : \mathcal{R}^{2n^2} \to \mathcal{R}^{n^2}$ has a $w_{n\times n}(u, v)$ Grigoriev's flow, where:*

$$w_{n\times n}(u, v) \geq \frac{1}{2}\left(v - \frac{\left(2n^2 - u\right)^2}{4n^2}\right), \text{ for } 0 \leq u \leq 2n^2,\ 0 \leq v \leq n^2. \tag{1}$$

The *dominator set* concept was originally introduced in [15].

▶ **Definition 3** (Dominator set)**.** Given a CDAG $G(V, E)$, let $I \subset V$ denote the set of input vertices. A *dominator set* for $V' \subseteq V$ is a set $\Gamma \subseteq V$ such that every path from a vertex in I to a vertex in V' contains at least a vertex of Γ. A *minimum dominator set* for V' is a dominator set with minimum cardinality.

We will use a specular concept, commonly referred as "*post-dominator set*".

▶ **Definition 4** (Post-dominator set)**.** Given a CDAG $G(V, E)$, let $O \subset V$ denote the set of output vertices. A *post-dominator set* for $V' \subseteq V \setminus O$ with respect to $O' \subseteq O$ is defined to be a set of vertices in V such that every path from a vertex in V' to the output vertices in O' contains at least a vertex of the set. A *minimum post-dominator* set for $V' \subseteq V \setminus O$ with respect to $O' \subseteq O$ is a post-dominator set with minimum cardinality.

The following lemma relates the size of certain post-dominator sets to the Grigoriev's flow.

▶ **Lemma 5.** *Let $G(V, E)$ be a CDAG computing $f : \mathcal{R}^p \to \mathcal{R}^q$ with Grigoriev's flow $w_f(u, v)$. Let I (resp., O) denote the set of input (resp., output) vertices of G. Any post-dominator set Γ for any subset $I' \subseteq I$ with respect to any subset $O' \subseteq O$ satisfies $|\Gamma| \geq w_f(|I'|, |O'|)$.*

Proof. Given $I' \subseteq I$ and $O' \subseteq O$, suppose the values of the input variables corresponding to vertices in $I \setminus I'$ to be fixed. Let Γ be a post-dominator set for I' with respect to O'.
(i) According to the definition of flow of f, there exists an assignment of the input variables corresponding to vertices in I' such that the output variables in O' can assume $|\mathcal{R}|^{w_f(|I'|,|O'|)}$

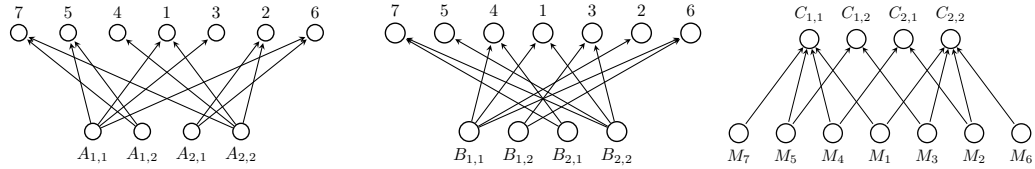

(a) Enc_A **(b)** Enc_B **(c)** Dec

Figure 1 Basic constructing blocks of Strassen's CDAG. Note that Enc_A and Enc_B are isomorphic.

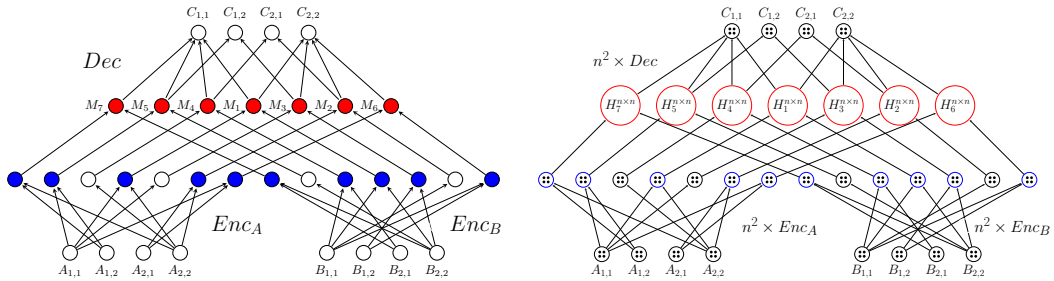
$C_{1,1}$ $C_{1,2}$ $C_{2,1}$ $C_{2,2}$
Dec
M_7 M_5 M_4 M_1 M_3 M_2 M_6
Enc_A
Enc_B
$A_{1,1}$ $A_{1,2}$ $A_{2,1}$ $A_{2,2}$
$B_{1,1}$ $B_{1,2}$ $B_{2,1}$ $B_{2,2}$
$n^2 \times Dec$
$H_7^{n\times n}$ $H_5^{n\times n}$ $H_1^{n\times n}$ $H_4^{n\times n}$ $H_3^{n\times n}$ $H_2^{n\times n}$ $H_6^{n\times n}$
$n^2 \times Enc_A$
$n^2 \times Enc_B$

(a) Strassen's $H^{2\times 2}$ CDAG **(b)** Recursive construction of the $H^{2n\times 2n}$ CDAG

Figure 2 Blue vertices represent combinations of the input values from the factor matrices A and B which are used as input values for the sub-problems M_i; red vertices represent the output of the seven sub-problems which are used to compute the output values of the product matrix C.

distinct values. (ii) As there is no path from I' to O' which has not a vertex in Γ, the values of the outputs in O' are determined by the inputs in $I \setminus I'$, which are fixed, and the values corresponding to the vertices in Γ; hence the outputs in O' can assume at most $|\mathcal{R}|^{|\Gamma|}$ distinct values. The claimed results follows by a simple combination of observations (i) and (ii). ◀

Properties of Strassen's Algorithm

Consider Strassen's algorithm [24] when used to compute $C = AB$, with A and B matrices $n \times n$ with entries from the ring $\mathcal{R}$ (see Algorithm 1 in Appendix A for a detailed presentation). Let $H^{n\times n}$ denote the corresponding CDAG. For $n \geq 2$, $H^{n\times n}$ can be obtained by using a recursive construction which mirrors the recursive structure of the algorithm. The base of the construction is the $H^{2\times 2}$ CDAG which corresponds to the multiplication of two 2×2 matrices using Strassen's algorithm (Figure 2a). $H^{2n\times 2n}$ can then be constructed by composing seven copies of $H^{n\times n}$, each corresponding to one of the seven sub-products generated by the algorithm (see Figure 2b): n^2 vertex-disjoint copies of CDAGs Enc_A (resp., Enc_B) are used to connect the input vertices of $H^{2n\times 2n}$ which correspond to the values of the input matrix A (resp., B) to the appropriate input vertices of the seven sub-CDAGs $H_i^{n\times n}$; the output vertices sub-CDAGs $H_i^{n\times n}$ (which correspond to outputs of the seven sub-products) are connected to the opportune output vertices of the entire $H^{2n\times 2n}$ CDAG using n^2 copies of the decoder sub-CDAG Dec.

In our proofs we will leverage the following two properties of Strassen's CDAG:

▶ **Lemma 6.** *Let $H^{n\times n}$ denote the CDAG of Strassen's algorithm for input matrices of size $n \times n$. For $0 \leq i \leq \log n - 1$, there are exactly 7^i sub-CDAGs $H^{n/2^i \times n/2^i}$ which do not share any vertex in $H^{n\times n}$ (i.e., they are* vertex disjoint sub-CDAGs of $H^{n\times n}$*).*

▶ **Lemma 7.** *Given an encoder CDAG, for any subset Y of its output vertices, there exists a sub-set X of its input vertices such that each vertex in X can be connected to a distinct vertex in Y where* $\min\{|Y|, 1 + \lceil (|Y| - 1)/2 \rceil\} \leq |X| \leq |Y|$.

The validity of this lemmas can be verified through an analysis of respectively, the recursive structure of Strassen's CDAG and the properties of the encoding sub-CDAGs Enc_A and Enc_B. We refer the reader to Appendix A for detailed proofs.

Strassen-like algorithms

A (n_0, m_0)-Strassen-like algorithm has a recursive structure for which in the "*base case*" two $n_0 \times n_0$ matrices are multiplied using m_0 scalar multiplications, whose result are then combined to obtain the product matrix. Given input matrices of size $n \times n$, the algorithm splits them into n_0^2 sub-matrices of size $n/n_0 \times n/n_0$ and then proceeds block-wise, according to the base case. Additions (resp., subtractions) in the base case are interpreted as additions (resp., subtractions) of blocks and are performed element-wise. Multiplications in the base case are interpreted as multiplications of blocks and are executed by recursively calling the algorithm. In our analysis we consider Strassen-like algorithms for which each linear combination of the input sub-matrices is used in only one multiplication.

Let $\mathcal{H}^{n\times n}$ be the CDAG corresponding to a (n_0, m_0)-Strassen-like algorithm for input matrices of size $n \times n$. $\mathcal{H}^{n\times n}$ has a recursive structure analogous to that of Strassen's CDAG $H^{n\times n}$. The base element $\mathcal{H}^{n_0\times n_0}$ of the CDAG construction corresponding to the base case consists of two *encoding graphs* (corresponding respectively to the sub-CDAGs Enc_A and Enc_B for $H^{2\times 2}$), which compute m_0 linear combinations of entries of respectively the factor matrix A and of B. Corresponding pairs are then multiplied and the outputs are then combined using an *decoder* sub-CDAG to obtain the output.

Although in general the sub-problems generated by a Strassen-like algorithm may not be input disjoint, in [23] Scott et al. showed that a property similar to Lemma 6 holds:

▶ **Lemma 8** (Lemma 1 [23])**.** *Let $\mathcal{H}^{n\times n}$ denote the CDAG of a (n_0, m_0)-Strassen-like algorithm for input matrices of size $n \times n$. For $0 \leq i \leq \log n - 1$, there are at least m_0^{i-2} vertex-disjoint sub-CDAGs $\mathcal{H}^{n/n_0^i \times n/n_0^i}$ in $\mathcal{H}^{n\times n}$.*

For completeness, we summarize their result in Appendix A.

3 Lower bounds for schedules without recomputation

In our presentation we denote as $G^{n\times n}$ the CDAG corresponding to the execution of an *unspecified* algorithm which implements the square matrix multiplication function $f_{n\times n} : \mathcal{R}^{2n^2} \to \mathcal{R}^{n^2}$. Let $G^{q,n\times n}$ be a CDAG composed by q vertex-disjoint CDAGs $G^{n\times n}$. The set I (resp., O) of input (resp., output) vertices of $G^{q,n\times n}$ is given by the union of the sets of the input (resp., output) vertices of the q sub-CDAGs $G^{n\times n}$. According to this definition, the CDAG of Strassen $H^{n\times n}$ is a $G^{7^i, n/2^i \times n/2^i}$ CDAG for $i \in 0, 1, \ldots, \log_2 n - 1$ (Lemma 6).

▶ **Lemma 9.** *Given $G^{q,n\times n}$, let $O' \in O$ be a subset of its the output vertices. For any subset Γ of the vertices of $G^{q,n\times n}$ with $|\Gamma| \leq 2|O'|$, the set $I' \subseteq I$ of the input vertices which are not post-dominated by Γ satisfies $|I'| \geq 2n\sqrt{|O'| - 2|\Gamma|}$.*

Proof. Let O'_j (resp., Γ_j) denote the subset of O' (resp., Γ) which correspond to vertices in the j-th sub CDAG $G_j^{n\times n}$ for $j \in \{1, 2, \ldots, q\}$. As, by hypothesis, the sub-CDAGs $G_q^{n\times n}$

are vertex disjoint, $O'_1, O'_2, \ldots, O'_q$ (resp., $\Gamma_1, \Gamma_2, \ldots, \Gamma_q$) constitute a partition of O' (resp., Γ).

Let $I''_j \subseteq \mathcal{X}$ denote the set of input vertices of $G_j^{n\times n}$ for which Γ_j is a post-dominator with respect to O_j. From Lemma 2 and Lemma 5 the following condition must hold:

$$|\Gamma_j| \geq w_{n\times n} \geq \frac{1}{2}\left(|O'_j| - \frac{\left(2n^2 - |I''_j|\right)^2}{4n^2}\right). \tag{2}$$

Let $I'_j = I_j \setminus I''_j$ denote the set of input vertices of $G_j^{n\times n}$ for which Γ_j is a post-dominator with respect to O_j. Since $|I| = 2n^2$, from (2) we have $|I'_j|^2 \geq 4n^2\left(|O'_j| - 2|\Gamma_j|\right)$.

As the sub-CDAGs $G_j^{n\times n}$ are vertex disjoint, we can state that all the vertices in $I' = \cup_{j=1}^{i} I'_j$ are not post-dominated by Γ and $|I'|^2 = \sum_{j=1}^{q} |I'_j|^2 \geq 4n^2 \sum_{j=1}^{q}\left(|O'_j| - 2|\Gamma_j|\right) = 4n^2\left(|O'| - 2|\Gamma|\right)$. ◀

▶ **Corollary 10.** *Given $G^{q,n\times n}$, the minimum size of a dominator set of any subset of the output vertices $O' \in O$ is at least $\lceil |O'|/2 \rceil$.*

These results allows us to obtain a general theorem on the I/O complexity of matrix multiplication algorithms under the no-recomputation assumption.

▶ **Theorem 11** (Lower bound I/O complexity matrix multiplication with no recomputation)**.** *Let $G^{\mathcal{A}}$ the CDAG corresponding to an algorithm $\mathcal{A}$ which computes the product of two square matrices $A, B \in \mathcal{R}^{n^2}$. Suppose that $G^{\mathcal{A}}$ has q vertex disjoint sub-CDAGs $G^{2\sqrt{M}\times 2\sqrt{M}}$ each of which corresponds to input-disjoint sub-products of matrices of size $2\sqrt{M} \times 2\sqrt{M}$ generated by $\mathcal{A}$. Assuming no intermediate result is ever computed more than once, the I/O -complexity of $\mathcal{A}$ when run on a sequential machine with a cache of size M is:*

$$IO_{G^{\mathcal{A}}}(M) \geq qM. \tag{3}$$

If run on P processor each equipped with a local memory of size M the I/O complexity is:

$$IO_{G^{\mathcal{A}}}(P, M) \geq qM/P. \tag{4}$$

Proof. As first observation, note that the q vertex-disjoint sub-CDAGs $G^{2\sqrt{M}\times 2\sqrt{M}\times n}$ constitute a $G^{q,2\sqrt{M}\times 2\sqrt{M}}$ CDAG as previously defined. The set X of input (resp. Y of output) vertices of $G^{q,2\sqrt{M}\times 2\sqrt{M}}$ is composed by the union of the input (resp., output) vertices of the q sub-CDAGs $G^{2\sqrt{M}\times 2\sqrt{M}}$. Clearly, $|X| = |Y| = q4M$.

We start by proving (3). Let $\mathcal{C}$ be any computational schedule for the sequential execution of the algorithm $\mathcal{A}$ using a cache of size M for which each intermediate value is computed just once. Each intermediate value must therefore be kept in memory (either cache or slow) until all operations using it as an operand have been computed. The operations corresponding to vertices in $G^{\mathcal{A}}$ are executed according to a topological ordering of the vertices, that is, any value corresponding to a vertex in $u \in X$ which can be connected through a directed path to a vertex in $v \in Y$ will be computed in $\mathcal{C}$ before v. We partition $\mathcal{C}$ into q segments such that during each segment $\mathcal{C}_j$ the values corresponding to exactly $4M$ vertices in Y (denoted as Y_j) are computed for the first (and only) time. We shall now verify that at least M I/O operations are to be executed during every segment $\mathcal{C}_j$. The proof is by contradiction. Let Γ_j denote the set of vertices of $G^{\mathcal{A}}$ which correspond to the at most M values which are stored in the cache at the beginning of the segment and the at most $M-1$ values which are loaded into the cache form the slow memory during $\mathcal{C}_j$ by means of a *read* I/O operation. We thus have $|\Gamma_j| \leq 2M - 1$. In order for the $4M$ values corresponding to vertices in Y_j

to be computed during the segment without any additional I/O operation, there must be no path connecting any vertex in Y_i to any vertex in I which does not have at least one vertex in Γ_j (i.e. Γ_j has to be a *dominator set* of Y_j). If any such path exists then there is a previously computed intermediate result v which is required for the computation of at lest one of the values corresponding to a vertex in Y_j which is neither residing in the cache at the beginning of $\mathcal{C}_j$, nor loaded from slow memory to the cache. As no intermediate result can be computed twice this would lead to a contradiction. From Corollary 10, we have that any sub-set of $4M$ elements of X has dominator size at least $2M$. This leads to a contradiction.

At least M I/O operations are thus executed during each segment $\mathcal{C}_j$. Since, by construction, the q segments are not overlapping, we can conclude that at least qM I/O operations are necessary for the execution of $\mathcal{C}$. This concludes the proof for lower bound in (3).

The proof for the lower bound for the parallel model in equation (4), follows a similar strategy. If P processors are being used, at least one such processor P^* must compute at least $|Y|/P = q4M/P$ values corresponding to vertices in Y. The bound follows by applying the same argument discussed for the sequential case to the computation executed by P^*. ◀

This general result can be steadily applied to Strassen-like algorithms.

▶ **Corollary 12.** *Consider Strassen's algorithm being used to multiply two square matrices $A, B \in \mathcal{R}^{n\times n}$. Assuming no intermediate result is ever computed more than once, the I/O-complexity of the algorithm run on a sequential machine with a cache of size M is:*

$$IO_{H^{n\times n}}(M) \geq \frac{1}{7}\left(\frac{n}{\sqrt{M}}\right)^{\log_2 7} M. \tag{5}$$

If run on P processor each equipped with local memory of size $M \leq n^2$ the I/O complexity is:

$$IO_{H^{n\times n}}(P, M) \geq \frac{1}{7}\left(\frac{n}{\sqrt{M}}\right)^{\log_2 7} \frac{M}{P}. \tag{6}$$

Additionally, the I/O-complexity of any (n_0, m_0)-Strassen-like algorithm run on a sequential machine with a cache of size M is:

$$IO_{\mathcal{H}^{n\times n}}(M) = \Omega\left(\left(\frac{n}{\sqrt{M}}\right)^{\log_{n_0} m_0} M\right). \tag{7}$$

If run on P processor each equipped with local memory of size $M \leq n^2$ the I/O complexity is:

$$IO_{\mathcal{H}^{n\times n}}(P, M) = \Omega\left(\left(\frac{n}{\sqrt{M}}\right)^{\log_{n_0} m_0} \frac{M}{P}\right). \tag{8}$$

Proof. We provide a simple proof for the sequential cases in (5) and (7). Let us assume without loss of generality that $n = 2^a$ and $\sqrt{M} = 2^b$ for some $a, b \in \mathbb{N}$. At least $3M$ I/O operations are necessary in order to read all the $2n^2$ input values form slow memory to the cache and to write the n^2 output values to the slow memory once they have been computed. The statement of the theorem is therefore trivially verified if $n \leq 2\sqrt{M}$. For $n \geq 2\sqrt{M}$, the result in (5) follows from applying Theorem 11 to the $H^{n\times n}$ CDAG which, from Lemma 6, has $\left(n/2\sqrt{M}\right)^{\log_2 7}$ vertex-disjoint sub CDAGs $H^{2\sqrt{M}\times 2\sqrt{M}}$. The result in (7) similarly follows from applying Lemma 8. The results for the parallel model in (6) and (8) can be obtained using a generalization similar to the one described in Theorem 11. ◀

While our lower bounds correspond asymptotically to the known bounds in [23], our technique yields a simpler analysis, especially for the I/O complexity of Strassen-like algorithms. Our proof technique is based on the analysis of the recursive structure of Strassen-like algorithms and on the identification of the vertex-disjoint sub-CDAGs corresponding to the various sub-problems. The fact that Theorem 11 applies for *any* square matrix multiplication algorithm, allows us to obtain significant bounds without a detailed analysis of the CDAG corresponding to the *specific* algorithm being considered.

This property further suggests that our technique may be amenable to deal with hybrid "*non-stationary*" multiplication algorithms, which allow mixing of schemes of the previous Strassen-like class in different recursive levels (for the "*uniform*" sub-class) [11], or even within the same level (for the "*non-uniform*" sub-class).

4 Lower bounds for schedules with recomputation

Under no-recomputation assumption once an input value is loaded in memory or an intermediate result is calculated it is then necessary to maintain it in memory (either cache or slow) until the result of each operation which uses it as an input argument has been evaluated. This constitutes the foundation of the lower bound technique discussed in Theorem 11 as well as several other techniques discussed in literature (among others, the "*dichotomy width* technique" [9], the "*boundary flow* technique" [20]), including those yielding I/O lower bound for Strassen's algorithm [2, 23]. If recomputation is allowed, intermediate results can instead be deleted from *all* memory and recomputed. This introduces a substantial complication in the theoretical analysis of the I/O cost (see [4] for an extensive discussion).

In this section we present a new lower bound technique which yields a novel asymptotically tight I/O lower bound for Strassen's algorithm both in the sequential and parallel model. We start by presenting some technical lemmas which will then be used for the proof of our main result in Theorem 15.

▶ **Lemma 13.** *Let $H^{n\times n}$ be the CDAG which corresponds to the execution of Strassen's matrix multiplication algorithm for input matrices $A, B \in \mathcal{R}^{n\times n}$, with $n \geq 2\sqrt{M}$. Let $\mathcal{Y}$ (resp., $\mathcal{Z}$) denote the set of input (resp., output) vertices of the $\left(n/2\sqrt{M}\right)^{\log_2 7}$ sub-CDAGs $H^{2\sqrt{M}\times 2\sqrt{M}}$. Furthermore, let Γ be a subset of the set of* internal *(i.e., not input) vertices of the sub-CDAGs $H^{2\sqrt{M}\times 2\sqrt{M}}$. Let $\mathcal{X}$ denote the set on input vertices of $H^{n\times n}$. For any subset $Z \subseteq \mathcal{Z}$ such that $|Z| \geq 2|\Gamma|$ there exists a set $X \subset \mathcal{X}$ with $|X| \geq 4\sqrt{M\,(|Z| - 2|\Gamma|)}$ such that vertices in X can be connected to vertices in a subset $Y \in \mathcal{Y}$, with $|X| = |Y|$ via vertex disjoint paths. Furthermore all vertices in Y can be connected to a vertex in Z by a directed path which does not include any vertex in Γ.*

Proof. The proof is by induction on the size of the input matrices.
Base: In the base case we have $n = 2\sqrt{M}$. We therefore have $H^{n\times n} = H^{2\sqrt{M}\times 2\sqrt{M}}$ and the sets $\mathcal{Y}$ and $\mathcal{X}$ coincide. We can verify that the statement holds by simply applying Lemma 9 as $H^{2\sqrt{M}\times 2\sqrt{M}}$ corresponds to a $G^{1,2\sqrt{M}\times 2\sqrt{M}}$ CDAG.
Inductive step: Let us now assume that the statement is verified for $H^{n\times n}$, with $n \geq 2\sqrt{M}$. We shall show that the statement is verified for $H^{2n\times 2n}$ as well. Let $H_1^{n\times n}, H_2^{n\times n}, \ldots, H_7^{n\times n}$ denote the seven sub-CDAGs of $H^{2n\times 2n}$, each corresponding to one of the seven sub-products generated by the first recursive step of Strassen's algorithm. Let Z_i (resp., $\mathcal{Y}_i$) denote the subset of Z (resp., $\mathcal{Y}$) which correspond to vertices in $H_i^{n\times n}$. As, from Lemma 6,

the seven sub-CDAGs $H_i^{n\times n}$ are vertex disjoint among themselves, $Z_1, Z_2, \ldots, Z_7$ (resp., $\mathcal{Y}_1, \mathcal{Y}_2, \ldots, \mathcal{Y}_7$) constitute a partition of Z (resp., $\mathcal{Y}$).

For $i \in \{1, 2, \ldots, 7\}$, let Γ_i be the subset of vertices in Γ in the sub-CDAGs $H_i^{n\times n}$. Again, as the seven sub-CDAGs $H_i^{n\times n}$ are vertex disjoint among themselves, we have that $\Gamma_1, \Gamma_2, \ldots, \Gamma_7$ is a partition of Γ. This implies $\sum_{i=1}^{7} |\Gamma_i| = |\Gamma|$. Let $\delta_i = \max\{0, |Z_i| - 2|\Gamma_i|\}$, we have $\delta = \sum_{i=1}^{7} \delta_i \geq |Z| - 2|\Gamma|$.

Applying the inductive hypothesis to each sub-CDAG $H_i^{n\times n}$, we have that there is a subset $Y_i \subseteq \mathcal{Y}_i$ with $|Y_i| \geq 4\sqrt{M\delta_i}$ such that vertices of Y_i are connected to vertices in Z_i via paths which do not include any vertex in Γ_i. Furthermore vertices in Y_i can be connected to a subset K_i of the input vertices of $H_i^{n\times n}$ with $|K_i| = |Y_i|$, using vertex-disjoint paths. Since the sub-CDAGs $H_i^{n\times n}$ are vertex disjoint, so are the paths connecting vertices in Y_i to vertices in K_i. In order to conclude our proof we need to show that is possible to extend at least $4\sqrt{M\left(|Z| - 2|\Gamma|\right)}$ of these paths to vertices in $\mathcal{X}$ while still being vertex disjoint.

According to the construction of Strassen's CDAG, vertices in $\mathcal{X}$ are connected to vertices in $K_1, K_2, \ldots, K_7$ by means of $2n^2$ encoding sub-CDAGs Enc_A and Enc_B (n^2 of each). For the construction of $H^{2n\times 2n}$, none of these encoding sub-CDAGs shares any input or output vertices. For any given encoder sub-CDAGs each of its output vertices belongs to a different sub-CDAG $H_i^{n\times n}$. This fact ensures that for a single sub-CDAG $H_i^{n\times n}$ it is possible to connect all the vertices in K_i to a subset of the vertices in $\mathcal{X}$ via vertex disjoint paths.

For each of the $2n^2$ encoder sub-CDAGs, let us consider the vector $\mathbf{y}_j \in \{0, 1\}^7$. We have that $\mathbf{y}_j[i] = 1$ if the corresponding i-th output vertex (respectively according to the numbering indicated in Figure 1a or Figure 1b) is in K_i or $\mathbf{y}_j[i] = 0$ otherwise. We therefore have that $|\mathbf{y}_j|$ corresponds to the number of output vertices of the j-th encoder sub-CDAG which are in K. From Lemma 7, we have that for each encoder sub-CDAG there exists a subset $X_j \in \mathcal{X}$ of the input vertices of the j-th encoder sub-CDAG for which is possible to connect each vertex in X_j to a distinct output vertex of the j-th encoder sub-CDAG using vertex disjoint paths, each constituted by a singular edge with $\min\{|\mathbf{y}_j|, 1+\lceil(|\mathbf{y}_j| - 1)/2\rceil\} \leq |X_j| \leq |\mathbf{y}_j|$. The number of vertex disjoint paths connecting vertices in $\mathcal{X}$, to vertices in $\cup_{i=1}^{7} K_i$ is therefore at least $\sum_{j=1}^{2n^2} \min\{|\mathbf{y}_j|, 1 + \lceil(|\mathbf{y}_j| - 1)/2\rceil\}$, under the constraint that $\sum_{j=1}^{2n^2} \mathbf{y}_j[i] = 4\sqrt{M\delta^i}$. Without loss of generality, let us assume that $\delta_1 \geq \delta_2 \geq \ldots \geq \delta_7$. As previously stated, it is possible to connect all vertices in K_1 to vertices in $\mathcal{X}$ through vertex disjoint paths. Consider now all possible dispositions of the vertices in $\cup_{i=2}^{7} K_i$ over the outputs of the $2n^2$ encoder sub-CDAGs. Recall that the output vertices of an encoder sub-CDAG belong each to a different $H^{n\times n}$ sub-CDAG. From Lemma 7, we have that for each encoder, there exists a subset $X_j \subset X$ of the input vertices of the j-th encoder sub-CDAG, with $|X_j| \geq \min\left\{|\mathbf{y}_j|, 1 + \lceil(|\mathbf{y}_j| - 1)/2\rceil\right\} \geq \mathbf{y}_j[1] + \left(\sum_{i=2}^{7} \mathbf{y}_j[i]\right)/2$, for which is possible to connect all vertices in X_j to $|X_j|$ *distinct* output vertices of the j-th encoder sub-CDAG which are in $\cup_{i=1}^{7} K_i$ using $|X_j|$, thus using vertex disjoint paths. As all the *Enc* sub-CDAGs are vertex disjoint, we can sum their contributions and we can therefore conclude that the number of vertex disjoint paths connecting values in $\mathcal{X}$ to vertices in $\cup_{i=1}^{7} K_i$ is at least $|K_1| + \frac{1}{2}\sum_{i=2}^{7} |K_i| = 4\sqrt{M}\left(\sqrt{\delta_1} + \frac{1}{2}\sum_{i=2}^{7}\sqrt{\delta_i}\right)$. Squaring this quantity leads to:

$$\left(4\sqrt{M}\left(\sqrt{\delta_1} + \frac{1}{2}\sum_{i=2}^{7}\sqrt{\delta_i}\right)\right)^2 = 16M\left(\delta_1 + \sqrt{\delta_1}\sum_{i=2}^{7}\sqrt{\delta_i} + \left(\frac{1}{2}\sum_{i=2}^{7}\sqrt{\delta_i}\right)^2\right).$$

As, by assumption, $\delta_1 \geq \delta_2 \geq \ldots \delta_7$, we have: $\sqrt{\delta_1}\sqrt{\delta_i} \geq \delta_i$ for $i = 2, 3, \ldots, 7$. Thus:

$$\left(4\sqrt{M}\left(\sqrt{\delta_1} + \frac{1}{2}\sum_{i=2}^{7}\sqrt{\delta_i}\right)\right)^2 \geq 16M\sum_{i=1}^{7}\delta_i \geq \left(4\sqrt{M\left(|Z| - 2|\Gamma|\right)}\right)^2.$$

There are therefore at least $4\sqrt{M\left(|Z| - 2|\Gamma|\right)}$ vertex disjoint paths connecting vertices in $\mathcal{X}$ to vertices in $\cup_{i=2}^{7} K_i$. The lemma follows. ◀

▶ **Lemma 14.** *Let $H^{n\times n}$ be the CDAG which corresponds to the execution of Strassen's matrix multiplication algorithm for input matrices $A, B \in \mathcal{R}^{n^2}$, with $n \geq 2\sqrt{M}$. Let $\mathcal{Y}$ (resp., $\mathcal{Z}$) denote the set of input (resp., output) vertices of the $\left(n/2\sqrt{M}\right)^{\log_2 7}$ sub-CDAGs $H^{2\sqrt{M}\times 2\sqrt{M}}$. Any dominator set for any subset $Z \subseteq \mathcal{Z}$, has size at least $|Z|/2 = 2M$.*

Proof. The proof is by contradiction. Let Γ be a dominator set for Z in $H^{n\times n}$ such that $|\Gamma| \leq 2M-1$. Let $\Gamma' \subseteq \Gamma$ be the subset of "internal" vertices of any of any of the sub-CDAGs $H^{2\sqrt{M}\times 2\sqrt{M}}$ in Γ. Let $\mathcal{X}$ denote the set the input vertices of $H^{n\times n}$. From Lemma 13 we have that there exist at least $4\sqrt{M\left(|Z| - 2|\Gamma'|\right)}$ paths connecting a subset $X \subseteq \mathcal{X}$ of global input vertices to vertices in Z which do not include any vertex in Γ'. Further, the sub-paths of such paths that connect vertices in X to vertices in $Y \subseteq \mathcal{Y}$ are vertex disjoint among themselves. This implies that each of the vertices in $\Gamma \setminus \Gamma'$ can therefore be traversed by at most one of these paths. The number of directed paths from X to Z which do not traverse any vertex in Γ will therefore be at least $4\sqrt{M\left(|Z| - 2|\Gamma'|\right)} - \left(|\Gamma| - |\Gamma'|\right)$. We have:

$$\begin{aligned}
&\left(4\sqrt{M\left(|Z| - 2|\Gamma'|\right)} - \left(|\Gamma| - |\Gamma'|\right)\right)^2 \\
&\qquad = 16M\left(|Z| - 2|\Gamma'|\right) + \left(|\Gamma| - |\Gamma'|\right)^2 - 8\sqrt{M\left(|Z| - 2|\Gamma'|\right)}\left(|\Gamma| - |\Gamma'|\right) \\
&\qquad \geq 16M\left(|Z| - 2|\Gamma'|\right) - 8\sqrt{M\left(|Z| - 2|\Gamma'|\right)\left(|\Gamma| - |\Gamma'|\right)} \\
&\qquad \geq 16M\left(|Z| - 2|\Gamma|\right) + 32M\left(|\Gamma| - |\Gamma'|\right) - 8\sqrt{M\left(|Z| - 2|\Gamma'|\right)}\left(|\Gamma| - |\Gamma'|\right) \\
&\qquad \geq 16M\left(|Z| - 2|\Gamma|\right) + \left(|\Gamma| - |\Gamma'|\right)\left(32M - 8\sqrt{M\left(|Z| - 2|\Gamma'|\right)}\right)
\end{aligned}$$

Since by hypothesis $|Z| - 2|\Gamma'| \leq 4M$, we have $\left(4\sqrt{M\left(|Z| - 2|\Gamma'|\right)} - \left(|\Gamma| - |\Gamma'|\right)\right)^2 \geq 16M\left(|Z| - 2|\Gamma|\right)$, and thus $4\sqrt{M\left(|Z| - 2|\Gamma'|\right)} - \left(|\Gamma| - |\Gamma'|\right) \geq 4\sqrt{M\left(|Z| - 2|\Gamma|\right)}$.

For $|\Gamma| \leq 2M - 1$ there is therefore at least one path connecting an input vertex to a vertex in Z which does not traverse any vertex in Γ. This constitutes a contradiction. ◀

Lemma 14 provides us with the tools required to obtain our main result.

▶ **Theorem 15** (Lower bound I/O complexity Strassen's algorithm)**.** *Consider Strassen's algorithm being used to multiply two square matrices $A, B \in \mathcal{R}^{n\times n}$. The I/O-complexity of Strassen's algorithm when run on a sequential machine with a cache of size M is:*

$$IO_{H^{n\times n}}(M) \geq \frac{1}{7}\left(\frac{n}{\sqrt{M}}\right)^{\log_2 7} M. \tag{9}$$

If run on P processors each equipped with a local memory of size M the I/O complexity is:

$$IO_{H^{n\times n}}(P, M) \geq \frac{1}{7}\left(\frac{n}{\sqrt{M}}\right)^{\log_2 7} \frac{M}{P} \tag{10}$$

Proof. We start by proving (3). We assume without loss of generality that $n = 2^a$ and $\sqrt{M} = 2^b$ for some $a, b \in \mathbb{N}$. At least $3M$ I/O operations are necessary in order to read all the $2n^2$ input values form slow memory to the cache and to write the n^2 output values to the slow memory once they have been computed. The theorem is therefore verified if $n \leq 2\sqrt{M}$.

For $n \geq 4\sqrt{M}$, let $\mathcal{Z}$ denote the set of output vertices of the $\left(n/2\sqrt{M}\right)^{\log_2 7}$ sub-CDAGs $H^{2\sqrt{M}\times 2\sqrt{M}}$ of $H^{n\times n}$. Let $\mathcal{C}$ be any computation schedule for the sequential execution of Strassen's algorithm using a cache memory of size M. We partition $\mathcal{C}$ into segments such that during each segment $\mathcal{C}_i$ the values corresponding to exactly $4M$ distinct vertices in $\mathcal{Z}$ (denoted as Z_i) are computed for the *first time*. As we have $|\mathcal{Z}| = 4M\left(\frac{n}{2\sqrt{M}}\right)^{\log 7}$, there will be $\left(n/\left(2\sqrt{M}\right)\right)^{\log 7}$ such segments. We shall now verify that at least M I/O operations are to be executed during every segment $\mathcal{C}_i$. The proof is by contradiction. Let Γ_i denote the set of vertices of $H^{n\times n}$ which correspond to the at most M values which are stored in the cache at the beginning of the i-th segment and to the at most $M-1$ values loaded into the cache form the slow memory during $\mathcal{C}_i$ by means of a *read* I/O operation. We then have $|\Gamma_i| \leq 2M - 1$. In order for the $4M$ values from Z_i to be computed during the segment without any additional I/O operation, there must be no path connecting any vertex in Z_i to any input vertex of $H^{n\times n}$ which does not have at least one vertex in Γ_i (i.e. Γ_i has to be a *dominator set* of Z_i). From Lemma 14, we have that any sub-set of $4M$ elements of $\mathcal{Z}$ has dominator size at least $2M$. This leads to a contradiction.

At least M I/O operations are thus executed during each segment $\mathcal{C}_i$. Since, by construction, the $\left(\frac{n}{2\sqrt{M}}\right)^{\log 7}$ segments are not overlapping, we can therefore conclude that at least $M\left(\frac{n}{2\sqrt{M}}\right)^{\log 7}$ I/O operations are necessary for the execution of any computation schedule $\mathcal{C}$. This concludes the proof for lower bound for the sequential case in (9).
The proof for the bound for the parallel model in equation (10), follows a similar strategy: at least one of the P processors being used, denoted as P^*, must compute at least $|\mathcal{Z}|/P = 4M\left(\frac{n}{2\sqrt{M}}\right)^{\log 7}/P$ values corresponding to vertices in $\mathcal{Z}$. The bound follows by applying the same argument discussed for the sequential case to the computation executed by P^*. ◀

In [1], Ballard et al. presented a version of Strassen's algorithm whose I/O cost matches, up to a constant the lower bound obtained in Theorem 15. This implies that our bound is indeed asymptotically tight and that the version of Strassen's algorithm presented by Ballard et al. in [1], is indeed asymptotically optimal with respect to the I/O cost. Furthermore, as in the optimal algorithm presented in [1] no intermediate result is ever recomputed, we can conclude the use of recomputation can lead at most to a constant factor reduction of the I/O complexity for the execution of Strassen's matrix multiplication algorithm.

5 Conclusions

This work has contributed to the characterization of the I/O complexity of Strassen's algorithm, by establishing asymptotically tight lower bounds that hold even when recomputation is allowed. The technique we have developed crucially exploits the recursive nature of the CDAG, which makes it promising for the analysis of other recursive algorithms, beginning with fast rectangular matrix multiplication algorithms [16].

The relationship we have exploited between dominator size and Grigoriev's flow points

at connections between I/O complexity, (pebbling) space-time tradeoffs [22], and VLSI area-time tradeoffs [25]; these connections deserve further attention.

Some CDAGs for which non trivial I/O complexity lower bounds are known only in the case of no recomputations are described in [9]. These CDAGs are of particular interest as examples of speedups superlinear with the number of processors, in the "*limiting technology*", defined by fundamental limitations on device size and message speed. Whether such speedups hold even when recomputation is allowed remains an open question, which the techniques introduced here might help answer.

In general, while it is well known that recomputation may reduce the I/O complexity of some CDAGs, we are far from a characterization of those CDAGs for which recomputation is effective. This broad goal remains a challenge for any attempt toward a general theory of the communication requirements of computations.

Acknowledgements. This work was supported, in part, by MIUR of Italy under project AMANDA 2012C4E3KT 004 and by the University of Padova under projects STPD08JA32, CPDA121378/12, and CPDA152255/15.

References

1 Grey Ballard, James Demmel, Olga Holtz, Benjamin Lipshitz, and Oded Schwartz. Communication-optimal parallel algorithm for strassen's matrix multiplication. In *Proceedings of the twenty-fourth annual ACM symposium on Parallelism in algorithms and architectures*, pages 193–204. ACM, 2012.

2 Grey Ballard, James Demmel, Olga Holtz, Benjamin Lipshitz, and Oded Schwartz. Graph expansion analysis for communication costs of fast rectangular matrix multiplication. In *Design and Analysis of Algorithms*, pages 13–36. Springer, 2012.

3 Grey Ballard, James Demmel, Olga Holtz, and Oded Schwartz. Communication-optimal parallel and sequential cholesky decomposition. *SIAM Journal on Scientific Computing*, 32(6):3495–3523, 2010.

4 Grey Ballard, James Demmel, Olga Holtz, and Oded Schwartz. Minimizing communication in numerical linear algebra. *SIAM Journal on Matrix Analysis and Applications*, 32(3):866–901, 2011.

5 Grey Ballard, James Demmel, Olga Holtz, and Oded Schwartz. Graph expansion and communication costs of fast matrix multiplication. *Journal of the ACM (JACM)*, 59(6):32, 2012.

6 Gianfranco Bilardi and Enoch Peserico. A characterization of temporal locality and its portability across memory hierarchies. In *Automata, Languages and Programming*, pages 128–139. Springer, 2001.

7 Gianfranco Bilardi, Andrea Pietracaprina, and Paolo D'Alberto. On the space and access complexity of computation dags. In *Graph-Theoretic Concepts in Computer Science*, pages 47–58. Springer, 2000.

8 Gianfranco Bilardi and Franco P Preparata. Horizons of parallel computation. *Journal of Parallel and Distributed Computing*, 27(2):172–182, 1995.

9 Gianfranco Bilardi and Franco P. Preparata. Processor-time trade offs under bounded speed message propagation: Part ii, lower bounds. *Theory of Computing Systems*, 32(5):531–559, 1999.

10 Lynn E Cannon. A cellular computer to implement the kalman filter algorithm. Technical report, DTIC Document, 1969.

11 Craig C Douglas, Michael Heroux, Gordon Slishman, and Roger M Smith. Gemmw: a portable level 3 blas winograd variant of strassen's matrix-matrix multiply algorithm. *Journal of Computational Physics*, 110(1):1–10, 1994.

12 Matteo Frigo, Charles E Leiserson, Harald Prokop, and Sridhar Ramachandran. Cache-oblivious algorithms. In *Foundations of Computer Science, 1999. 40th Annual Symposium on*, pages 285–297. IEEE, 1999.

13 Dmitrii Yur'evich Grigor'ev. Application of separability and independence notions for proving lower bounds of circuit complexity. *Zapiski Nauchnykh Seminarov POMI*, 60:38–48, 1976.

14 Dror Irony, Sivan Toledo, and Alexander Tiskin. Communication lower bounds for distributed-memory matrix multiplication. *Journal of Parallel and Distributed Computing*, 64(9):1017–1026, 2004.

15 Hong Jia-Wei and Hsiang-Tsung Kung. I/o complexity: The red-blue pebble game. In *Proceedings of the thirteenth annual ACM symposium on Theory of computing*, pages 326–333. ACM, 1981.

16 François Le Gall. Faster algorithms for rectangular matrix multiplication. In *Foundations of Computer Science (FOCS), 2012 IEEE 53rd Annual Symposium on*, pages 514–523. IEEE, 2012.

17 Francois Le Gall. Powers of tensors and fast matrix multiplication. In *Proceedings of the 39th International Symposium on Symbolic and Algebraic Computation (ISSAC 2014)*, pages 296–303. ACM, 2014.

18 L. H. Loomis and H. Whitney. An inequality related to the isoperimetric inequality. *Bull. Amer. Math. Soc.*, 55(10):961–962, 10 1949.

19 Cynthia A Patterson, Marc Snir, Susan L Graham, et al. *Getting Up to Speed:: The Future of Supercomputing*. National Academies Press, 2005.

20 Desh Ranjan, John Savage, and Mohammad Zubair. Upper and lower i/o bounds for pebbling r-pyramids. *Journal of Discrete Algorithms*, 14:2–12, 2012.

21 John E Savage. Extending the hong-kung model to memory hierarchies. In *Computing and Combinatorics*, pages 270–281. Springer, 1995.

22 John E. Savage. *Models of Computation: Exploring the Power of Computing*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1st edition, 1997.

23 Jacob Scott, Olga Holtz, and Oded Schwartz. Matrix multiplication i/o complexity by path routing. In *Proceedings of the 27th ACM on Symposium on Parallelism in Algorithms and Architectures*, pages 35–45. ACM, 2015.

24 Volker Strassen. Gaussian elimination is not optimal. *Numerische Mathematik*, 13(4):354–356, 1969.

25 Clark David Thompson. Area-time complexity for vlsi. In *Proceedings of the eleventh annual ACM symposium on Theory of computing*, pages 81–88. ACM, 1979.

26 Yu. D. Burago V. A. Zalgaller, A. B. Sossinsky. *The American Mathematical Monthly*, 96(6):544–546, 1989.

27 Shmuel Winograd. On multiplication of 2× 2 matrices. *Linear algebra and its applications*, 4(4):381–388, 1971.

A Properties of Strassen's algorithm

Algorithm 1 Strassen's Matrix Multiplication

Input: matrices A, B
Output: matrix C
procedure StrassenMM(A,B)
  if $n = 1$ **then**
    $C = A \cdot B$
  else
    Decompose A and B into four equally sized bloc matrices as follows:

$$A = \begin{bmatrix} A_{1,1} & A_{1,2} \\ A_{2,1} & A_{2,2} \end{bmatrix}, \quad B = \begin{bmatrix} B_{1,1} & B_{1,2} \\ B_{2,1} & B_{2,2} \end{bmatrix}$$

    $M_1 = \text{StrassenMM}\,(A_{1,1} + A_{2,2}, B_{1,1} + B_{2,2})$
    $M_2 = \text{StrassenMM}\,(A_{2,1} + A_{2,2}, B_{1,1})$
    $M_3 = \text{StrassenMM}\,(A_{1,1}, B_{1,2} - B_{2,2})$
    $M_4 = \text{StrassenMM}\,(A_{2,2}, B_{2,1} - B_{1,1})$
    $M_5 = \text{StrassenMM}\,(A_{1,1} + A_{1,2}, B_{2,2})$
    $M_6 = \text{StrassenMM}\,(A_{2,1} - A_{1,1}, B_{1,1} + B_{1,2})$
    $M_7 = \text{StrassenMM}\,(A_{1,2} - A_{2,2}, B_{2,1} + B_{2,2})$
    $C_{1,1} = M_1 + M_4 - M_5 + M_7$
    $C_{1,2} = M_3 + M_5$
    $C_{2,1} = M_2 + M_4$
    $C_{2,2} = M_1 - M_2 + M_3 + M_6$
  return C

The original version of Strassen's fast matrix multiplication [24] is reported in Algortihm 1. We refer the reader to [27] for Winograd's variant, which reduces the number of additions.

Proof of Lemma 6. Let us consider a single recursive step. While some of the sub-problems may use as input values form either A and B without any combination (i.e., M_3 uses the sub-matrix $A_{1,1}$ as input), none of the seven sub-problems share any of their input values. As this consideration holds for each level of the recursive construction, we have that the 7^i sub-problems with input of size $\frac{n}{2^i} \times \frac{n}{2^i}$ generated by of Strassen's algorithm do not share any input value. ◀

Proof of Lemma 7. We provide the proof for Enc_A, as Enc_A and Enc_B are isomorphic, the result holds for Enc_B as well. Note that in Enc_A there are some pairs of input-output vertices u, v for which the input vertex u is the only predecessor of the output vertex v. This implies that the two vertices are actually one unique vertex. With a little abuse of notation, we will still say that u can be connected to v via a single edge.

We assign an index to each of the output vertices of Enc_A according to how is indicated in Figure 1a. Note that the index assigned to each output corresponds to the index of the sub-problem generated by Strassen's algorithm for which the corresponding value is used as an input (see Figure 2a and Figure 2b in Section 2).

In order to verify that this lemma holds, we study all possible compositions of a subset Y of the output vertices of Enc_A. Each of these compositions is identified by a vector $\mathbf{y}$ with

seven components, where $y_i = 1$ if the i-th output of Enc_A is in Y or zero otherwise, for $i \in \{1, 2, \ldots, 7\}$. In order to improve the presentation, we associate to each of the possible 128 compositions a *code* given by $\sum_{i=1}^{7} y_i 2^{7-i}$. In Table 1, we study each of the 128 possible compositions of Y, which are ordered by the value of $|Y|$ and by their code. The value in the last column $|X|$ denotes the maximum size of a sub-set X of the input vertices of Enc_A such that each vertex in X can be connected to a distinct vertex in the subset Y corresponding to $\mathbf{y}$. Each of these values can be obtained through a straightforward analysis of Enc_A. The lemma then follows from observing that for any possible composition of the set Y we have

$$|X| \geq \min\left\{|Y|, 1 + \lceil (|Y| - 1)/2 \rceil\right\}.$$

◀

Proof of Lemma 8. A proof for this result was presented in [23] (Lemma 1). For completeness, we present here a version of the proof using our notation.

Given an (n_0, m_0)-Strassen like algorithm used to multiply $A, B \in \mathcal{R}^{n^2}$, let $\mathcal{H}^{n \times n}$ denote the corresponding CDAG. The recursive structure of the algorithm will generate a total of $m_0^{\log_{n_0} n - i}$ sub-problems with input size n/n_0^i each of which corresponds to a sub-CDAGs $\mathcal{H}^{n/n_0^i \times n/n_0^i}$.

The lemma is trivially verified for $0 \leq i \leq 1$, in the following we consider the case for $2 \leq i \leq \log_{n_0} m_0$. As by hypothesis, non-trivial combinations of the input matrices are used as input for at most one sub-problem, the only input values which two sub-problems of the same size can share are sub-matrices of the "*global*" input matrices A, B. For $i' = i - 2$, et $\mathcal{P}_1$ be one of the sub-computations corresponding to one of the sub-CDAGs $\mathcal{H}^{n/n_0^{i'} \times n/n_0^{i'}}$ which multiplies matrices A_1 and B_1. Without loss of generality we assume that for both the basic encoder CDAGs for A and B at least one of the outputs is given by a non-trivial linear combination of elements of the input matrix. It is in fact well known that any algorithm which computes linear combinations of only one of the input matrices performs no better than the naíve-definition based matrix multiplication and it is therefore not a Strassen-like fast matrix multiplication algorithm. This implies that at least one of the sub-problems $\mathcal{P}_2$ generated by $\mathcal{P}_1$ multiplies matrices A_2 and B_2 such that A_2 is a nontrivial linear combination of sub-matrices of A_1. Similarly, at least one of the sub-computations $\mathcal{P}_3$ generated by $\mathcal{P}_2$ multiplies matrices A_3 by B_3 such that B_3 is non-trivial linear combination of sub-matrices of B. A_3 (resp., B_3) shares no input with A (resp., B). Thus at least one of the m_0^2 sub-computations for input size $n/n_0^i \times n/n_0^i$ generated by $\mathcal{P}_1$ is input-disjoint from it. As by hypothesis, non-trivial combinations of the input matrices are used as input for at most one sub-problem, all the sub-computations $\mathcal{P}_3$ corresponding to each $\mathcal{P}_1$ are input-disjoint and the corresponding sub-CDAGs $\mathcal{H}^{n/n_0^i \times n/n_0^i}$ are thus vertex disjoint. ◀

Table 1 Study of the possible compositions of sub-sets of output vertices of *Enc* for Lemma 7

code	$\lvert\mathbf{y}\rvert$	y_1	y_2	y_3	y_4	y_5	y_6	y_7	$\lvert X\rvert$
0	0	0	0	0	0	0	0	0	0
1	1	0	0	0	0	0	0	1	1
2	1	0	0	0	0	0	1	0	1
4	1	0	0	0	0	1	0	0	1
8	1	0	0	0	1	0	0	0	1
16	1	0	0	1	0	0	0	0	1

Continued on next page

Table 1 - continued from previous page

code	$\|\mathbf{y}\|$	y_1	y_2	y_3	y_4	y_5	y_6	y_7	$\|X\|$
32	1	0	1	0	0	0	0	0	1
64	1	1	0	0	0	0	0	0	1
3	2	0	0	0	0	0	1	1	2
5	2	0	0	0	0	1	0	1	2
6	2	0	0	0	0	1	1	0	2
9	2	0	0	0	1	0	0	1	2
10	2	0	0	0	1	0	1	0	2
12	2	0	0	0	1	1	0	0	2
17	2	0	0	1	0	0	0	1	2
18	2	0	0	1	0	0	1	0	2
20	2	0	0	1	0	1	0	0	2
24	2	0	0	1	1	0	0	0	2
33	2	0	1	0	0	0	0	1	2
34	2	0	1	0	0	0	1	0	2
36	2	0	1	0	0	1	0	0	2
40	2	0	1	0	1	0	0	0	2
48	2	0	1	1	0	0	0	0	2
65	2	1	0	0	0	0	0	1	2
66	2	1	0	0	0	0	1	0	2
68	2	1	0	0	0	1	0	0	2
72	2	1	0	0	1	0	0	0	2
80	2	1	0	1	0	0	0	0	2
96	2	1	1	0	0	0	0	0	2
7	3	0	0	0	0	1	1	1	3
11	3	0	0	0	1	0	1	1	3
13	3	0	0	0	1	1	0	1	3
14	3	0	0	0	1	1	1	0	3
19	3	0	0	1	0	0	1	1	3
21	3	0	0	1	0	1	0	1	3
22	3	0	0	1	0	1	1	0	3
25	3	0	0	1	1	0	0	1	3
26	3	0	0	1	1	0	1	0	3
28	3	0	0	1	1	1	0	0	3
35	3	0	1	0	0	0	1	1	3
37	3	0	1	0	0	1	0	1	3
38	3	0	1	0	0	1	1	0	3
41	3	0	1	0	1	0	0	1	3
42	3	0	1	0	1	0	1	0	3
44	3	0	1	0	1	1	0	0	3
49	3	0	1	1	0	0	0	1	3
50	3	0	1	1	0	0	1	0	3
52	3	0	1	1	0	1	0	0	3
56	3	0	1	1	1	0	0	0	3
67	3	1	0	0	0	0	1	1	3
69	3	1	0	0	0	1	0	1	3

Continued on next page

Table 1 - continued from previous page

code	$\|\mathbf{y}\|$	y_1	y_2	y_3	y_4	y_5	y_6	y_7	$\|X\|$
70	3	1	0	0	0	1	1	0	3
73	3	1	0	0	1	0	0	1	3
74	3	1	0	0	1	0	1	0	3
76	3	1	0	0	1	1	0	0	3
81	3	1	0	1	0	0	0	1	3
82	3	1	0	1	0	0	1	0	3
84	3	1	0	1	0	1	0	0	3
88	3	1	0	1	1	0	0	0	2
97	3	1	1	0	0	0	0	1	3
98	3	1	1	0	0	0	1	0	3
100	3	1	1	0	0	1	0	0	3
104	3	1	1	0	1	0	0	0	3
112	3	1	1	1	0	0	0	0	3
15	4	0	0	0	1	1	1	1	4
23	4	0	0	1	0	1	1	1	4
27	4	0	0	1	1	0	1	1	4
29	4	0	0	1	1	1	0	1	3
30	4	0	0	1	1	1	1	0	4
39	4	0	1	0	0	1	1	1	4
43	4	0	1	0	1	0	1	1	4
45	4	0	1	0	1	1	0	1	4
46	4	0	1	0	1	1	1	0	4
51	4	0	1	1	0	0	1	1	4
53	4	0	1	1	0	1	0	1	4
54	4	0	1	1	0	1	1	0	4
57	4	0	1	1	1	0	0	1	4
58	4	0	1	1	1	0	1	0	3
60	4	0	1	1	1	1	0	0	4
71	4	1	0	0	0	1	1	1	4
75	4	1	0	0	1	0	1	0	3
77	4	1	0	0	1	1	0	1	3
78	4	1	0	0	1	1	1	0	4
83	4	1	0	1	0	0	1	1	4
85	4	1	0	1	0	1	0	1	3
86	4	1	0	1	0	1	1	0	4
89	4	1	0	1	1	0	0	1	3
90	4	1	0	1	1	0	1	0	3
92	4	1	0	1	1	1	0	0	3
99	4	1	1	0	0	0	1	1	4
101	4	1	1	0	0	1	0	1	4
102	4	1	1	0	0	1	1	0	4
105	4	1	1	0	1	0	0	1	4
106	4	1	1	0	1	0	1	0	3
108	4	1	1	0	1	1	0	0	4
113	4	1	1	1	0	0	0	1	4

Continued on next page

Table 1 - continued from previous page

code	$\|\mathbf{y}\|$	y_1	y_2	y_3	y_4	y_5	y_6	y_7	$\|X\|$
114	4	1	1	1	0	0	1	0	3
116	4	1	1	1	0	1	0	0	4
120	4	1	1	1	1	0	0	0	3
31	5	0	0	1	1	1	1	1	4
47	5	0	1	0	1	1	1	1	4
55	5	0	1	1	0	1	1	1	4
59	5	0	1	1	1	0	1	1	4
61	5	0	1	1	1	1	0	1	4
62	5	0	1	1	1	1	1	0	4
79	5	1	0	0	1	1	1	1	4
87	5	1	0	1	0	1	1	1	4
91	5	1	0	1	1	0	1	1	4
93	5	1	0	1	1	1	0	1	3
94	5	1	0	1	1	1	1	0	4
103	5	1	1	0	0	1	1	1	4
107	5	1	1	0	1	0	1	1	4
109	5	1	1	0	1	1	0	1	4
110	5	1	1	0	1	1	1	0	4
115	5	1	1	1	0	0	1	1	4
117	5	1	1	1	0	1	0	1	4
118	5	1	1	1	0	1	1	0	4
121	5	1	1	1	1	0	0	1	4
122	5	1	1	1	1	0	1	0	3
124	5	1	1	1	1	1	0	0	4
63	6	0	1	1	1	1	1	1	4
95	6	1	0	1	1	1	1	1	4
111	6	1	1	0	1	1	1	1	4
119	6	1	1	1	0	1	1	1	4
123	6	1	1	1	1	0	1	1	4
125	6	1	1	1	1	1	0	1	4
126	6	1	1	1	1	1	1	0	4
127	7	1	1	1	1	1	1	1	4